

filterVcf: Extract Variants of Interest from a Large VCF File

Paul Shannon

June 21, 2015

Contents

1	Introduction	2
2	The Data: Paired Tumor/Normal Breast Cancer Variants	2
3	Filter by Genomic Region	3
4	Introducing the filterVcf Method	3
4.1	Prefilters	4
4.2	Filters	4
4.3	FilterRules	5
4.4	Create the Filtered file	5
5	Look for SNPs in Regulatory Regions	6
5.1	Load CTCF Transcription Factor Binding Regions Identified in MCF-7 Breast Cancer Cell Line	6
5.2	Find SNPs in CTCF Binding Regions	7
6	Conclusion	7

1 Introduction

Whole genome Variant Call Format (VCF) files are very large, typically containing millions of called variants, one call per line of text. The actual number of variants relevant to a particular study (of a disease such as breast cancer, for instance) will often be far fewer. Thus the first task one faces when analyzing a whole genome VCF file is to identify and extract the relatively few variants which may be of interest, excluding all others. This vignette illustrates several techniques for doing this.

We demonstrate three methods: filtering by genomic region, filtering on attributes of each specific variant call, and intersecting with known regions of interest (exons, splice sites, regulatory regions, etc.). We are primarily concerned with the latter two. However, in order to create the small VCF data file we use here for demonstration purposes, we employed genomic region filtering, reducing a very large whole genome two-sample breast cancer VCF file of fourteen million calls, to a file containing fewer than ten thousand calls in a one million base pair region of chromosome 7. For the sake of reproducibility, and for completeness of exposition, we will illustrate this first step also.

2 The Data: Paired Tumor/Normal Breast Cancer Variants

Complete Genomics Inc.¹ states:

To provide the scientific community with public access to data generated from two paired tumor/normal cancer samples, Complete Genomics sequenced and analyzed cell-line samples of patients with breast cancer (invasive ductal carcinomas). The cell line-derived DNA are housed at ATCC. Samples have been sequenced to an average genome-wide coverage of 123X for three of the samples, and 92X for for the fourth sample.²

¹<http://www.completegenomics.com/public-data/cancer-data>

²Data generated with version 2.0.0.32 of the Complete Genomics assembly software, on high molecular weight genomic DNA isolated from the HCC1187 breast carcinoma cell line ATCC CRL 2322. Sequencing methods documented in [Drmanac et al, 2010].

A small (1M base) subset of this data is included in the current package, and used in the code presented below.

3 Filter by Genomic Region

We identified this subset in a prior exploration of the full data set (work not shown), learning that variants of biological interest suitable for our purpose are found in a 1M base pair region of chromosome seven. The appendix to this document (see below) shows the few lines of code required to extract variant calls in that small region, from the very large file obtained from Complete Genomics ("*somaticVcfBeta-HCC1187-H-200-37-ASM-T1-N1.vcf.gz*") and write them to a new, small VCF file.

4 Introducing the filterVcf Method

The `filterVcf` method reads (by chunks, about which more below) through a possibly very large VCF file to write a new, smaller VCF file which includes only those variant calling rows which meet the criteria specified in `prefilters` and `filters`.

Reading "by chunks" is accomplished using a tabix [Li (2011)] file. The `yieldSize` argument specifies how many variant lines are read into memory at each iteration.

```
tabix.file <- TabixFile(file.gz, yieldSize=10000)
filterVcf(tabix.file, genome, destination.file,
          prefilters=prefilters, filters=filters)
```

in which

1. *file.gz*: a gzipped vcf file with an accompanying Tabix index file.
2. *yieldSize*: the number of text (call variant) lines to read at a time.
3. *genome*: a string indicating the genome assembly, e.g., "hg19".
4. *prefilters*: one or more simple string-based filtering functions, each of which returns a logical vector corresponding to the vcf rows it will be passed (as simple character strings).
5. *filters*: one or more filtering function, each of which returns a logical vector, corresponding to the list of parsed vcf structures it will be passed.

4.1 Prefilters

Prefilters are conceptually very simple. They are functions which will be called with a single argument, a vector of character strings, each of which is an *unparsed* variant call line, as read from the input VCF file. We use `grepl` to return a logical vector equal in length to the incoming vector of unparsed VCF lines. Each prefilter and filter is called repeatedly, with `yieldSize` lines supplied on each invocation. `filterVcf` calls these functions repeatedly until the input file is exhausted.

Notice how the logic of these prefilters is very simple, using `grepl` to do fast, simple, fixed pattern matching:

```
isGermlinePrefilter <- function(x) {
  grepl("Germline", x, fixed=TRUE)
}
notInDbsnpPrefilter <- function(x) {
  !(grepl("dbsnp", x, fixed=TRUE))
}
```

4.2 Filters

Filters are more sophisticated than prefilters in that they assess *parsed* variant call lines for possible inclusion. Such parsing is intrinsically expensive but will be performed only on those lines which passed the prefilters. Therefore it pays to eliminate as many lines as possible using prefilters. Filters are useful when there exists detailed criteria for inclusion and exclusion. This can be seen below, especially in the `allelicDepth` function.

```
isSnp <- function(x) {
  refSnp <- nchar(ref(x)) == 1L
  a <- alt(x)
  altSnp <- elementLengths(a) == 1L
  ai <- unlist(a[altSnp]) # all length 1, so unlisting is 1:1 map
  altSnp[altSnp] <- nchar(ai) == 1L & (ai %in% c("A", "C", "G", "T"))
  refSnp & altSnp
}
allelicDepth <- function(x) {
  ## ratio of AD of the "alternate allele" for the tumor sample
  ## OR "reference allele" for normal samples to total reads for
  ## the sample should be greater than some threshold (say 0.1,
  ## that is: at least 10% of the sample should have the allele
  ## of interest)
```

```
ad <- geno(x)$AD
tumorPct <- ad[,1,2,drop=FALSE] / rowSums(ad[,1,,drop=FALSE])
normPct <- ad[,2,1, drop=FALSE] / rowSums(ad[,2,,drop=FALSE])
test <- (tumorPct > 0.1) | (normPct > 0.1)
!is.na(test) & test
}
```

4.3 FilterRules

`FilterRules` allow you to combine a list of filters, or of prefilters so that they may be passed as parameters to *filterVcf*. We use them here to combine the *isGermlinePrefilter* with the *notInDbsnpPrefilter*, and the *isSnp* with the *AD* filter.

```
library(VariantAnnotation)
prefilters <- FilterRules(list(germline=isGermlinePrefilter,
                               dbsnp=notInDbsnpPrefilter))
filters <- FilterRules(list(isSnp=isSnp, AD=allelicDepth))
```

4.4 Create the Filtered file

```
file.gz      <- system.file("extdata", "chr7-sub.vcf.gz",
                             package="VariantAnnotation")
file.gz.tbi  <- system.file("extdata", "chr7-sub.vcf.gz.tbi",
                             package="VariantAnnotation")
destination.file <- tempfile()
tabix.file <- TabixFile(file.gz, yieldSize=10000)
filterVcf(tabix.file, "hg19", destination.file,
           prefilters=prefilters, filters=filters, verbose=TRUE)
```

```

found header lines for 3 "fixed" fields: ALT, QUAL, FILTER
found header lines for 16 "info" fields: CGA_BF, CGA_BNDG, CGA_BNDGO, CGA_FI, CGA_MEDEL, CGA_MN, CGA_MN2, CGA_MN3, CGA_MN4, CGA_MN5, CGA_MN6, CGA_MN7, CGA_MN8, CGA_MN9, CGA_MN10, CGA_MN11, CGA_MN12, CGA_MN13, CGA_MN14, CGA_MN15, CGA_MN16, CGA_MN17, CGA_MN18, CGA_MN19, CGA_MN20, CGA_MN21, CGA_MN22, CGA_MN23, CGA_MN24, CGA_MN25, CGA_MN26, CGA_MN27, CGA_MN28, CGA_MN29, CGA_MN30, CGA_MN31, CGA_MN32, CGA_MN33, CGA_MN34, CGA_MN35, CGA_MN36, CGA_MN37, CGA_MN38, CGA_MN39, CGA_MN40, CGA_MN41, CGA_MN42, CGA_MN43, CGA_MN44, CGA_MN45, CGA_MN46, CGA_MN47, CGA_MN48, CGA_MN49, CGA_MN50, CGA_MN51, CGA_MN52, CGA_MN53, CGA_MN54, CGA_MN55, CGA_MN56, CGA_MN57, CGA_MN58, CGA_MN59, CGA_MN60, CGA_MN61, CGA_MN62, CGA_MN63, CGA_MN64, CGA_MN65, CGA_MN66, CGA_MN67, CGA_MN68, CGA_MN69, CGA_MN70, CGA_MN71, CGA_MN72, CGA_MN73, CGA_MN74, CGA_MN75, CGA_MN76, CGA_MN77, CGA_MN78, CGA_MN79, CGA_MN80, CGA_MN81, CGA_MN82, CGA_MN83, CGA_MN84, CGA_MN85, CGA_MN86, CGA_MN87, CGA_MN88, CGA_MN89, CGA_MN90, CGA_MN91, CGA_MN92, CGA_MN93, CGA_MN94, CGA_MN95, CGA_MN96, CGA_MN97, CGA_MN98, CGA_MN99, CGA_MN100, CGA_MN101, CGA_MN102, CGA_MN103, CGA_MN104, CGA_MN105, CGA_MN106, CGA_MN107, CGA_MN108, CGA_MN109, CGA_MN110, CGA_MN111, CGA_MN112, CGA_MN113, CGA_MN114, CGA_MN115, CGA_MN116, CGA_MN117, CGA_MN118, CGA_MN119, CGA_MN120, CGA_MN121, CGA_MN122, CGA_MN123, CGA_MN124, CGA_MN125, CGA_MN126, CGA_MN127, CGA_MN128, CGA_MN129, CGA_MN130, CGA_MN131, CGA_MN132, CGA_MN133, CGA_MN134, CGA_MN135, CGA_MN136, CGA_MN137, CGA_MN138, CGA_MN139, CGA_MN140, CGA_MN141, CGA_MN142, CGA_MN143, CGA_MN144, CGA_MN145, CGA_MN146, CGA_MN147, CGA_MN148, CGA_MN149, CGA_MN150, CGA_MN151, CGA_MN152, CGA_MN153, CGA_MN154, CGA_MN155, CGA_MN156, CGA_MN157, CGA_MN158, CGA_MN159, CGA_MN160, CGA_MN161, CGA_MN162, CGA_MN163, CGA_MN164, CGA_MN165, CGA_MN166, CGA_MN167, CGA_MN168, CGA_MN169, CGA_MN170, CGA_MN171, CGA_MN172, CGA_MN173, CGA_MN174, CGA_MN175, CGA_MN176, CGA_MN177, CGA_MN178, CGA_MN179, CGA_MN180, CGA_MN181, CGA_MN182, CGA_MN183, CGA_MN184, CGA_MN185, CGA_MN186, CGA_MN187, CGA_MN188, CGA_MN189, CGA_MN190, CGA_MN191, CGA_MN192, CGA_MN193, CGA_MN194, CGA_MN195, CGA_MN196, CGA_MN197, CGA_MN198, CGA_MN199, CGA_MN200, CGA_MN201, CGA_MN202, CGA_MN203, CGA_MN204, CGA_MN205, CGA_MN206, CGA_MN207, CGA_MN208, CGA_MN209, CGA_MN210, CGA_MN211, CGA_MN212, CGA_MN213, CGA_MN214, CGA_MN215, CGA_MN216, CGA_MN217, CGA_MN218, CGA_MN219, CGA_MN220, CGA_MN221, CGA_MN222, CGA_MN223, CGA_MN224, CGA_MN225, CGA_MN226, CGA_MN227, CGA_MN228, CGA_MN229, CGA_MN230, CGA_MN231, CGA_MN232, CGA_MN233, CGA_MN234, CGA_MN235, CGA_MN236, CGA_MN237, CGA_MN238, CGA_MN239, CGA_MN240, CGA_MN241, CGA_MN242, CGA_MN243, CGA_MN244, CGA_MN245, CGA_MN246, CGA_MN247, CGA_MN248, CGA_MN249, CGA_MN250, CGA_MN251, CGA_MN252, CGA_MN253, CGA_MN254, CGA_MN255, CGA_MN256, CGA_MN257, CGA_MN258, CGA_MN259, CGA_MN260, CGA_MN261, CGA_MN262, CGA_MN263, CGA_MN264, CGA_MN265, CGA_MN266, CGA_MN267, CGA_MN268, CGA_MN269, CGA_MN270, CGA_MN271, CGA_MN272, CGA_MN273, CGA_MN274, CGA_MN275, CGA_MN276, CGA_MN277, CGA_MN278, CGA_MN279, CGA_MN280, CGA_MN281, CGA_MN282, CGA_MN283, CGA_MN284, CGA_MN285, CGA_MN286, CGA_MN287, CGA_MN288, CGA_MN289, CGA_MN290, CGA_MN291, CGA_MN292, CGA_MN293, CGA_MN294, CGA_MN295, CGA_MN296, CGA_MN297, CGA_MN298, CGA_MN299, CGA_MN300, CGA_MN301, CGA_MN302, CGA_MN303, CGA_MN304, CGA_MN305, CGA_MN306, CGA_MN307, CGA_MN308, CGA_MN309, CGA_MN310, CGA_MN311, CGA_MN312, CGA_MN313, CGA_MN314, CGA_MN315, CGA_MN316, CGA_MN317, CGA_MN318, CGA_MN319, CGA_MN320, CGA_MN321, CGA_MN322, CGA_MN323, CGA_MN324, CGA_MN325, CGA_MN326, CGA_MN327, CGA_MN328, CGA_MN329, CGA_MN330, CGA_MN331, CGA_MN332, CGA_MN333, CGA_MN334, CGA_MN335, CGA_MN336, CGA_MN337, CGA_MN338, CGA_MN339, CGA_MN340, CGA_MN341, CGA_MN342, CGA_MN343, CGA_MN344, CGA_MN345, CGA_MN346, CGA_MN347, CGA_MN348, CGA_MN349, CGA_MN350, CGA_MN351, CGA_MN352, CGA_MN353, CGA_MN354, CGA_MN355, CGA_MN356, CGA_MN357, CGA_MN358, CGA_MN359, CGA_MN360, CGA_MN361, CGA_MN362, CGA_MN363, CGA_MN364, CGA_MN365, CGA_MN366, CGA_MN367, CGA_MN368, CGA_MN369, CGA_MN370, CGA_MN371, CGA_MN372, CGA_MN373, CGA_MN374, CGA_MN375, CGA_MN376, CGA_MN377, CGA_MN378, CGA_MN379, CGA_MN380, CGA_MN381, CGA_MN382, CGA_MN383, CGA_MN384, CGA_MN385, CGA_MN386, CGA_MN387, CGA_MN388, CGA_MN389, CGA_MN390, CGA_MN391, CGA_MN392, CGA_MN393, CGA_MN394, CGA_MN395, CGA_MN396, CGA_MN397, CGA_MN398, CGA_MN399, CGA_MN400, CGA_MN401, CGA_MN402, CGA_MN403, CGA_MN404, CGA_MN405, CGA_MN406, CGA_MN407, CGA_MN408, CGA_MN409, CGA_MN410, CGA_MN411, CGA_MN412, CGA_MN413, CGA_MN414, CGA_MN415, CGA_MN416, CGA_MN417, CGA_MN418, CGA_MN419, CGA_MN420, CGA_MN421, CGA_MN422, CGA_MN423, CGA_MN424, CGA_MN425, CGA_MN426, CGA_MN427, CGA_MN428, CGA_MN429, CGA_MN430, CGA_MN431, CGA_MN432, CGA_MN433, CGA_MN434, CGA_MN435, CGA_MN436, CGA_MN437, CGA_MN438, CGA_MN439, CGA_MN440, CGA_MN441, CGA_MN442, CGA_MN443, CGA_MN444, CGA_MN445, CGA_MN446, CGA_MN447, CGA_MN448, CGA_MN449, CGA_MN450, CGA_MN451, CGA_MN452, CGA_MN453, CGA_MN454, CGA_MN455, CGA_MN456, CGA_MN457, CGA_MN458, CGA_MN459, CGA_MN460, CGA_MN461, CGA_MN462, CGA_MN463, CGA_MN464, CGA_MN465, CGA_MN466, CGA_MN467, CGA_MN468, CGA_MN469, CGA_MN470, CGA_MN471, CGA_MN472, CGA_MN473, CGA_MN474, CGA_MN475, CGA_MN476, CGA_MN477, CGA_MN478, CGA_MN479, CGA_MN480, CGA_MN481, CGA_MN482, CGA_MN483, CGA_MN484, CGA_MN485, CGA_MN486, CGA_MN487, CGA_MN488, CGA_MN489, CGA_MN490, CGA_MN491, CGA_MN492, CGA_MN493, CGA_MN494, CGA_MN495, CGA_MN496, CGA_MN497, CGA_MN498, CGA_MN499, CGA_MN500, CGA_MN501, CGA_MN502, CGA_MN503, CGA_MN504, CGA_MN505, CGA_MN506, CGA_MN507, CGA_MN508, CGA_MN509, CGA_MN510, CGA_MN511, CGA_MN512, CGA_MN513, CGA_MN514, CGA_MN515, CGA_MN516, CGA_MN517, CGA_MN518, CGA
```

5 Look for SNPs in Regulatory Regions

We have now created a file containing 29 novel, Germline, SNP variant calls, each with a reasonable allelic depth, extracted from the 3808 calls in *chr7-sub.vcf*. We examine those variants for overlap with regulatory regions, turning to the ENCODE project and using Bioconductor's [AnnotationHub](#).

The ENCODE project is a large, long-term effort to build a “parts list” of all the functional elements in the human genome. They have recently focused on regulatory elements. We use the [BioconductorAnnotationHub](#) to download regulatory regions reported for a breast cancer cell line, with which to identify possibly functional, and possibly clinically relevant SNVs in the breast cancer tumor/normal genome we have been examining.

The [AnnotationHub](#) is a recent addition to Bioconductor. See here [TODO] for more information.

5.1 Load CTCF Transcription Factor Binding Regions Identified in MCF-7 Breast Cancer Cell Line

The *MCF-7* Breast Cancer Cell line (<http://en.wikipedia.org/wiki/MCF-7>) was established forty years ago, and has since played a dominant role in breast cancer cell line studies. The University of Washington reports transcription factor binding sites (TFBS) for the CTCF protein, which often acts as a negative regulator of transcription via chromatin structure modifications. Though the MCF-7 cell line is an imperfect match to the HCC1187 cell line sequenced by Complete Genomics, we combine these two breast-cancer related data sets here, didactically, in this exercise, to highlight the importance of cell-type-specific regulatory regions, and of the availability of such data from ENCODE. We shall see a SNP in the intronic binding region of the cancer-related gene, EGFR.

```
library(AnnotationHub)
hub <- AnnotationHub()
# paste operation is only for display purposes on a narrow page
ctcf.regs <- paste("goldenpath.hg19.encodeDCC.wgEncodeUwTfbs",
                  "wgEncodeUwTfbsMcf7CtcfStdPkRep1.narrowPeak_0.0.1.RData",
                  sep=".")
mcf7.gr <- hub[[ctcf.regs]]
```

5.2 Find SNPs in CTCF Binding Regions

```
vcf <- readVcf(destination.file, "hg19")
seqlevels(vcf) <- paste("chr", seqlevels(vcf), sep="")
ov.mcf7 <- findOverlaps(vcf, mcf7.gr)
```

There is just one SNV which overlaps with the MCF-7 regulatory regions. Find out where, if anywhere, it fits within a gene model.

```
library(TxDb.Hsapiens.UCSC.hg19.knownGene)
txdb <- TxDb.Hsapiens.UCSC.hg19.knownGene
print(locateVariants(vcf[6,], txdb, AllVariants()))
```

6 Conclusion

This case study begins, somewhat artificially, with a very short region of chromosome seven, a section which we knew, from previous exploration, held an intronic regulatory SNV for EGFR, a receptor tyrosine kinase implicated in some cancers. Though artificial, the case study illustrates all of the steps needed for broader, realistic surveys of whole genome variation data:

1. Filter by genomic coordinates.
2. Filter to extract only those variant calls which meet these criteria: Germline, novel (not in dbSnp), consisting of a single nucleotide, of sufficient allelic depth.
3. Intersect these variants with recognized DNA elements. In our case, we used short TFBS regulatory regions, but the same method can be used with exons, splice sites, DNaseI footprints, methylation sites, etc.

7 Appendix: Filter by Genomic Region

The most basic form of VCF file filtering is by genomic region. We demonstrate that here, extracting variant calls in a 1M base region of chromosome 7, writing them to a new file, compressing and then indexing that file. These steps created the small VCF file which accompanies this vignette, and is used in the code shown above.

Note that this code is *NOT* executed during the creation of this vignette: we do not supply the very large VCF file that it operates on. This code is here for tutorial purposes

only, showing how you can filter by genomic region with a possibly large VCF file of your own.

```
library(VariantAnnotation)
file.gz <- "somaticVcfBeta-HCC1187-H-200-37-ASM-T1-N1.vcf.gz"
stopifnot(file.exists(file.gz))
file.gz.tbi <- paste(file.gz, ".tbi", sep="")
if(!(file.exists(file.gz.tbi)))
  indexTabix(file.gz, format="vcf")
start.loc <- 55000000
end.loc <- 56000000
chr7.gr <- GRanges("7", IRanges(start.loc, end.loc))
params <- ScanVcfParam(which=chr7.gr)
vcf <- readVcf(TabixFile(file.gz), "hg19", params)
writeVcf(vcf, "chr7-sub.vcf")
bgzip("chr7-sub.vcf", overwrite=TRUE)
indexTabix("chr7-sub.vcf.gz", format="vcf")
```

This code creates the small gzipped vcf and index files used in the examples above.

8 Bibliography

Radoje Drmanac, Andrew B Sparks, Matthew J Callow, Aaron L Halpern, Norman L Burns, Bahram G Kermani, Paolo Carnevali, Igor Nazarenko, Geoffrey B Nilsen, George Yeung, et al. Human genome sequencing using unchained base reads on self-assembling dna nanoarrays. *Science*, 327(5961):78–81, 2010.

Heng Li. Tabix: fast retrieval of sequence features from generic tab-delimited files. *Bioinformatics*, 27(5):718–719, 2011.